Session / Séance 28-A

Area-Normalized Thematic Views

T. Alan Keahey

Los Alamos National Laboratory
keahey@lanl.gov

Abstract

Thematic variables are commonly used to encode additional information such as population density within the spatial layout of a map. Such «themes» are typically encoded using colour-maps. We will explore techniques for using this thematic information to directly define spatial transformations so that areas of map regions are proportional to their thematic variables, thus making the view more consistent with the thematic encodings. Our method emphasizes interactivity as a primary mechanism for allowing the user to better realize the distribution of the thematic variable, rather than relying on static views of the map.

Introduction

A common problem in cartography occurs when a thematic variable is used to control the shading of regions of the map; this can very often lead to a situation where the areas used to represent the various thematic values are not consistent with the values themselves, possibly leading the viewer to misinterpret the information that the map designer is trying to convey [Tufte, 1983]. As a simple example of this, imagine two countries having identical populations, but one has a geographic area one tenth the size of the other. On a thematic map showing population, the larger country will make a more significant visual impression on the viewer than the smaller one, despite their identical population values, possibly leading to the viewer attaching a greater weight to the larger country. Many efforts at reducing or eliminating this problem by distorting the map to make the areas of regions proportional to their thematic values, while still maintaining connectivity between regions, have been described in the cartography literature, a common term for such a method of solution is the *continuous cartogram*.

A different but related problem to this occurs in computer visualization where zooming in on regions of the display causes the viewer to lose his or her awareness of context, so that it is not possible to see both the big picture and the fine details at the same time (bringing to mind the saying “he couldn’t see the forest for the trees”). The concept of a *fisheye view* is often applied in such cases, so that the viewer can focus in on magnified regions of interest while still maintaining a sense of the overall context at reduced resolutions (thus allowing the viewer to see both the trees *and* the forest simultaneously). The simplest example of such a view is familiar to anyone who has looked through a fisheye lens of a camera or a security peephole installed in a door.

This paper will explore the regions of intersection between these two methods – continuous cartograms and fisheye views – and show how the tools and methods obtained from each of them can be applied to the other. We will begin with a generalized description of fisheye views (nonlinear magnification), followed by a more expressive abstraction for fisheye views (nonlinear magnification fields). This new abstraction can be applied to many of the same problems that continuous cartograms are attempting to address. We will emphasize the importance of interactivity for the user, and show that our system is efficient enough to produce results at near-interactive rates. We finish with a discussion of related work, conclusions, and plans for further work.

Nonlinear Magnification

Many approaches have been described in the literature for stretching and distorting spaces to produce effective visualizations. Examples include *fisheye views*, *stretching a rubber-sheet*, *focus+context*, and *distortion-oriented displays*. The term *nonlinear magnification* was introduced in [Keahey and Robertson, 1996] to describe the effects common to all of these approaches. The basic properties of nonlinear magnification are non-occluding in-place magnification which preserves a view of the global context. A brief discussion of some of the characteristics of more traditional systems for generating nonlinear magnification will provide some context and motivation for the section that follows. More details about these types of systems are available from the Nonlinear Magnification Home Page at *www.cs.indiana.edu/hyplan/tkeahey/research/nlm/*.

Traditionally, most nonlinear magnification systems begin with defining a center of magnification, often referred to as a *focus*. This center of magnification can be either a point, a line, or a region, depending on the specific system under discussion. The idea is that areas near this center of magnification will be enlarged, while the surrounding *context* areas will be compressed. The simplest examples of this type of magnification are illustrated in Figure 1, which shows both a radial fisheye-type transformation, as well as an orthogonal transformation where the horizontal and vertical axes are treated independently. These examples show the effect of a transformation applied to a regular grid of points, and the user is able to interactively change the location of the center of magnification as desired.

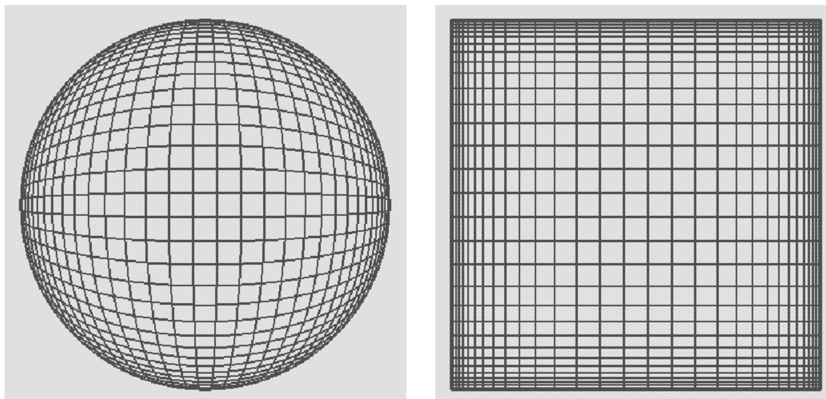

Figure 1. Radial and Orthogonal Transformations

From this simple starting point, many more complex transformations have been developed. Some examples of the kinds of transformations that are possible include: using regions of distortion-free linear magnification within the fisheye view, placing boundaries around the regions of magnification, and combining multiple centers of magnification in various ways. Examples of these are shown in Figure 2.

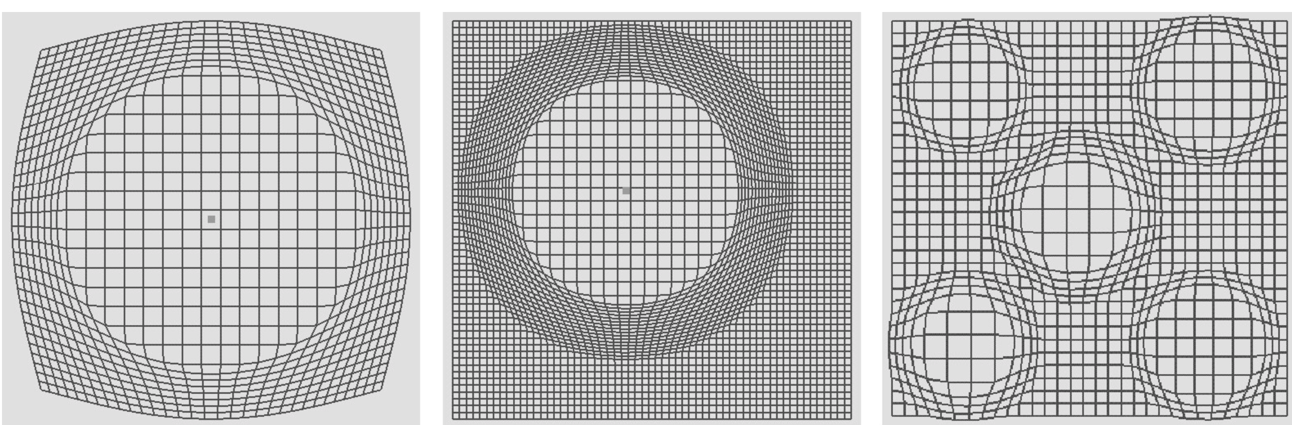

Figure 2. **Examples of More Complex Transformations**

The examples shown so far have just involved transformations on a regular grid of points, however there are many examples of applying nonlinear magnification to more irregular data types such as trees [Munzner, 1997], graphs, arbitrary polygons and GIS data [Churcher et al., 1997]. When we use a regular grid of points, we can use *texture mapping* [Blinn and Newell, 1976] to map any image (texture) onto that grid, so that as transformations are applied to the grid they are also applied to the mapped image. This makes transformations on a regular grid a particularly useful technique, as we can now easily transform arbitrary images at interactive rates using the texture mapping hardware acceleration that is now increasingly common on even low-end PC machines. An example of applying transformations to a texture map is seen in Figure 3, where we magnify a region on a map of the Washington D.C. area.

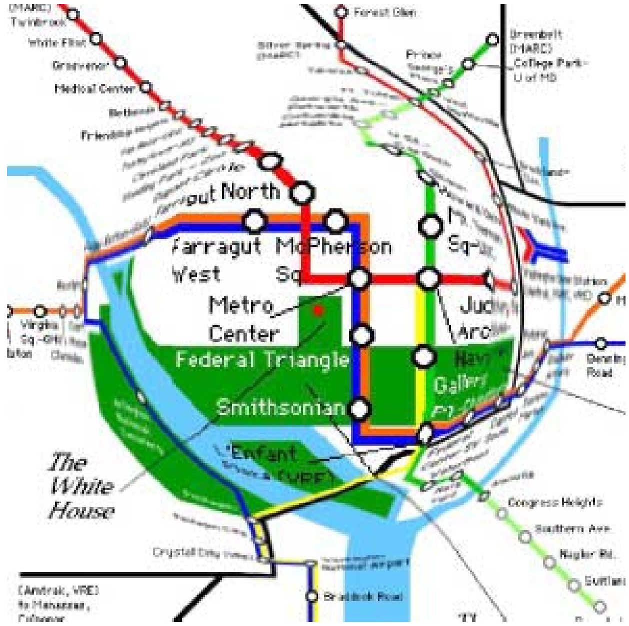

Figure 3. Using Texture Mapping to Transform an Image

Most of the more recent research efforts into nonlinear magnification have stressed the need for providing a high degree of interactivity for the user. User studies have emphasized that smooth operation and feedback of the magnification mechanism is necessary if the user is not to become disoriented [Schaffer et al., 1993]. These requirements typically can be met if the system is able to transform and display the information at a rate of 10 frames per second or greater. Some of the nonlinear magnification research efforts that describe an emphasis on computational efficiency are able to sustain interactive rates while transforming tens of thousands point coordinates for each frame [Keahey and Robertson, 1996; Munzner, 1997], thus illustrating the rapid transformation rates that are possible with these “foci-based” systems. We would like to maintain this high degree of interactivity for our thematic map transformations, as will be discussed in a later section.

There is some limitation on the degree of transformational expressiveness that is possible when using the foci-based systems however. Despite the many ways in which these centers of magnification can be constructed and combined, we are still conceptually limited to the idea of having discrete centers of magnification. Developing more complex transformations with such systems involves the addition of more foci, and it becomes a non-trivial matter to analytically predict the overall effect of these multiple interdependent foci [Keahey and Robertson, 1997]. An early effort at applying this type of foci-based magnification to the problem of continuous cartograms can be found in the polyfocal projection work of [Kadmon and Shlomi, 1978]. We will see in the next section how removing the restrictions of discrete foci can allow for a much more expressive class of transformations which is better suited for the complex transformations required of continuous cartograms.

Nonlinear Magnification Fields

Leung and Apperley [1994] first established the mathematical relationship between 1D magnification and transformation functions for nonlinear magnification systems, defining magnification as the derivative of the transformation function. This idea was extended to higher dimensions using an area-based derivative, resulting in the abstraction of a scalar field of magnification values called a *nonlinear magnification field* [Keahey and Robertson, 1997; Keahey, 1997]. As a result of this work, it is possible to compute the *implicit magnification field* of a given nonlinear transformation, which gives the magnification values inherent in the transformation. An example is shown in Figure 4, where a nonlinear transformation is shown beside its implicit magnification field.

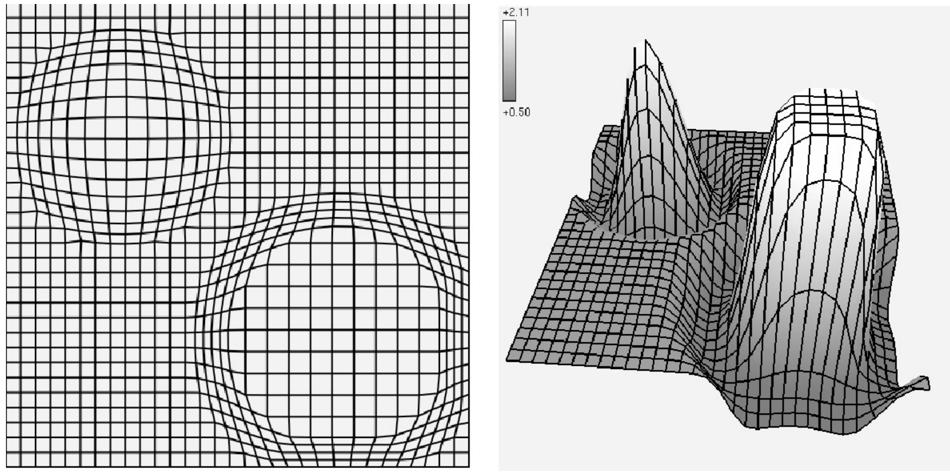

Figure 4. Transformation and Implicit Magnification Field

In addition, an iterative method is described in [Keahey and Robertson, 1997] that computes suitable spatial transformations based on a specified scalar field of magnification values. The scalar magnification field is particularly amenable to user and program manipulation, and provides a much more expressive class of transformations than is possible with traditional foci-based approaches to nonlinear magnification such as [Kadmon and Shlomi, 1978; Keahey and Robertson, 1996]. Briefly, the iterative method begins with an untransformed "working" grid and a mesh of desired magnification values. For each iteration, the method computes the implicit magnification field of the working grid and then subtracts the obtained magnification values from the desired magnification values to obtain a mesh of error values. Then for each node in the working grid if the associated error is positive we push away the nearest neighbours, and if the error is negative we pull the neighbors closer to that node. Full details of the iterative method are available in [Keahey, 1997]. With these magnification fields, it is now possible to directly define magnification values without having to deal with the side-effects of multi-dimensional transformation functions, and without concern for the complex interactions of multiple foci. The additional expressiveness of nonlinear magnification fields is crucial to the methods we present here; it is now possible to create *data-driven magnifications* [Keahey and Robertson, 1997], where properties of the data are used to *directly* define the magnification best suited for viewing that data.

Thematic Magnification

There is a natural match between data-driven magnification and the colour-encoding of thematic variables in maps. We can easily define routines which place a regular grid over a raster image of RGB values, and use the sampled RGB values to derive suitable magnification levels at each point in the grid, producing a magnification mesh. Given that magnification mesh, we can then compute a suitable transformation having the desired cartogram-like properties. The accuracy of this method is dependent on two major factors: the resolution of the map image (the degree to which its pixels represent an accurate sampling of the original data), and the resolution of the magnification mesh (the degree to which it accurately samples the pixels in the map image). In each case, sampling theory such as the Nyquist theorem gives guidelines as to what resolution would be required to accurately sample some given minimal signal frequency. Since our end goal here is only to achieve a visual approximation to the correct results, it is usually appropriate to relax the sampling requirements so that they are tuned more to the average-sized frequency components rather than to the smallest ones. For some extreme cases however, the sampling frequency requirements can be somewhat stringent, and may require the use of established multi-resolution methods for dealing with them most effectively.

Once the magnification mesh has been computed, a number of preprocessing operations can be performed on its values before the transformation is computed. Smoothing of the values with a low-pass filter is sometimes useful for cases with extreme variation from node to node (high-frequency spikes). Nonlinear scaling of the values can be used to emphasize particular ranges of magnification values. Other possibilities include locking individual nodes in the mesh, and constraining them to lie along certain vectors. During the iterative process, we can also weight the method to emphasize correction of undersized regions, oversize regions, or those regions which are maximally different from their desired size. Details of these operations are provided in [Keahey, 1997]. Complex effects can also be achieved by encoding different information in each RGB channel; the examples in this paper use the R channel to define the magnification values, and the G channel to specify logical "don't care" values for those areas of the map where the R values are not well defined (e.g. in the bodies of water surrounding geographic regions). This new ability to define "don't cares" is particularly advantageous in terms of computational efficiency. Whereas previous cartogram systems typically assigned some constant (average) value to these regions, in our iterative method we can simply ignore them and allow neighbouring regions to push and pull them without constraint (other than preserving mesh topology). Conversely, the "don't care" regions will not exert any influence on their neighbours. This can greatly reduce the amount of computation required for convergence.

Interaction Issues with Thematic Maps

Although there have been a number of methods for continuous cartograms which do a reasonably good job of producing a transformation of the map having the desired areas [Gusein-Zade and Tikunov, 1993], the problem of determining how well such cartograms convey the desired information to the viewer remains difficult. Recent work has effectively addressed part of this problem by including region recognizability as part of the iterative transformation process [House and Kocmoud, 1998]. Some fundamental issues on the effectiveness of cartograms are worthy of discussion at this point however.

A key difficulty with the conceptual notion of cartograms is that the human visual system is not very finely tuned for detecting slight differences in area between regions. Two examples of this are shown in Figure 5. Looking at the two images on the left, and trying to visually determine which of them has a larger area, you will find that without the use of measuring devices it is very difficult to say with certainty that the two objects have the same area. This example illustrates that the *shape* of a region can influence our estimate of its size; more complex shapes such as those found in cartograms will present even greater difficulties for the viewer. Next, examine the two images on the right in Figure 5, and try to visually determine if the central circle is larger in the left or right image. For most viewers, the circle on the left will initially appear to be larger, despite the fact that they are the same size. This example is the well-known Titchener illusion, which illustrates how even when the size and shape of an object remain constant, the surrounding context for that object can influence our perception of its size.

Because of the above difficulties, our method employs a shift of emphasis away from demanding *absolute* exact area representations, and towards creating perceptually recognizable transformations that provide the user with a better understanding of the *relative* distribution of the thematic variable. A key ingredient of our method is the use of animation and interactivity. Animation provides a means to smoothly interpolate between the regular and transformed views of the space, allowing the viewer to realize the relationship between the normal familiar view and the view which more accurately reflects the thematic content. Placing that animation under user control allows the viewer to rewind, playback, and pause, thus allowing the viewer to manipulate the independent thematic variable and obtain a "feel" for its distribution. These methods provide a direct binding between the complex independent thematic variable, and the simple variable of time. Thus the time variable provides a stable frame of reference for the user as he or she manipulates the complex variable to understand its properties.

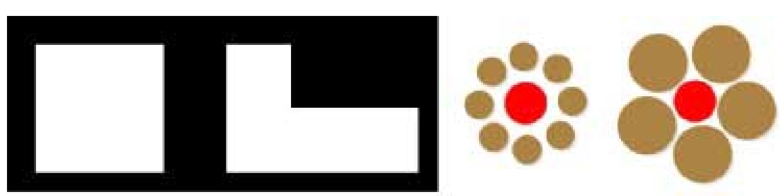

Figure 5. The Difficulty of Visually Comparing the Areas of Regions

When we observe a static cartogram only as a view at the end of a transformation process, we do not inherently have a description of the nature of the transformation that has taken place, excepting for those cases where the untransformed view of the data is so familiar to us that it serves as an internal source for comparison. For the less familiar cases, we can place a view of the untransformed map alongside the cartogram, this will provide some basis for comparison and realizing the distribution of the thematic variable, however subtle differences may still be lost or erroneously inferred as the viewer shifts her gaze between images. Through the use of animation and interactivity however, the user is able to smoothly compare the normal and transformed views without having to shift gaze, thus allowing much more subtle effects of the transformation to be recognized. In addition, the user can now easily track the progress of the transformation on a region that is familiar in the untransformed view, even if the final transformation produces a significantly distorted representation of that region.

In support of our requirements for animation and interactivity, the method that we use for computing the transformations performs at near-interactive rates. While other systems for continuous cartogram production can take many hours to compute [House and Kocmoud, 1998], ours can typically be computed in about the same amount of time it takes to read the map image from disk (from a fraction of a second to a few seconds, depending on the specific task). Animation between normal and transformed views can be achieved either through simple interpolation between the original and transformed regular grid, or through viewing time steps from the iterative process as it converges.

Example I: Presidential Election Results

The presidential election in the United States is decided by the number of electoral votes each candidate receives. Each state has a given number of electoral votes (based on the state population), and all of the electoral votes for a single state must be given entirely to only one of the candidates. It is common practice on election day for the news organizations to show a map of the USA, shading a state in blue (or dark gray) if they voted for the Democratic candidate, and red (or light gray) if they voted for the Republican candidate. This gives rise to a classic problem in information visualization that occurs when the area used to visually represent each region is not consistent with the actual thematic variable of importance [Tufte, 1983].

Figure 6 shows a traditional view of the presidential election results from 1996. If this image were to accurately reflect the number of electoral votes each candidate received we would expect the ratio of red (light) to blue (dark) pixels to be 0.42; what we actually get however is a ratio of 1.23, a global error of 193% which could leave the viewer to mistakenly infer that the Republican candidate (Dole) won the election instead of the Democratic candidate (Clinton). The error occurs because large and sparsely populated states such as Alaska and Montana visually dominate the image even though they have very few electoral votes, while states with a large number of electoral votes such as New York, Texas and California are not represented with an area-emphasis proportional to their electoral contributions. Note that the error measure used here is a global one based on the sum total of pixels in the image. This is in contrast to the per-region error measures often used in other cartogram systems such as [Dougenik et al., 1985; Gusein-Zade and Tikunov, 1993; House and Kocmoud, 1998], where the average of errors for the regions is used rather than the sum total of them.

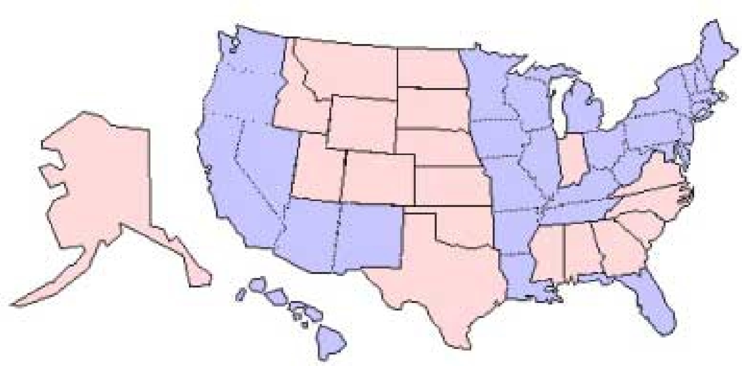

Figure 6. Traditional View of Election Results

To reduce this error we can construct a map of the USA where shading is used to represent the number of electoral votes in each state, as shown in the left image of Figure 7. The right image of that figure shows how the shading is used to define a magnification field for the theme.

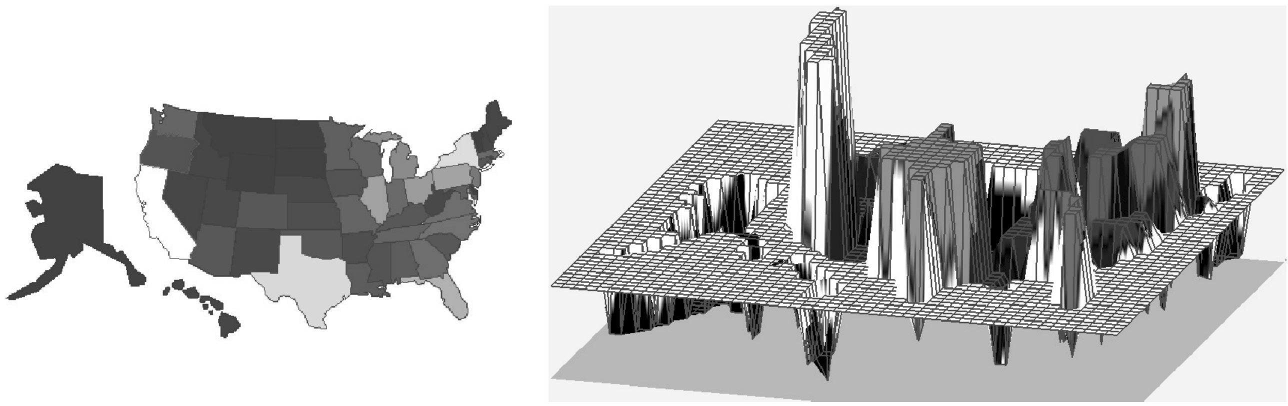

Figure 7. Shading by Electoral Votes and Construction of Magnification Field

We can then compute a magnification based on that thematic content to transform the normal view of the election into one that more accurately represents the actual proportion of electoral votes received by each candidate. The result is shown in Figure 8, where the ratio of red (light) to blue (dark) pixels is 0.69. Although this still represents a global error of 64%, this is less than 1/3 of the global error found in the original image, and the ratio of pixels now accurately reflects the fact that Clinton won the election. Since our transformation applies only to images and there is no representation of discrete state boundaries in our data, comparisons of this error with the averages of per-state errors described in [Dougenik et al., 1985; Gusein-Zade and Tikunov, 1993; House and Kocmoud, 1998] are difficult. One possibility would be to divide the total error by the number of states (50) to obtain an average error of 1.28% for each state, which would compare quite favourably with the previously mentioned systems. In reality however, there is probably some cancellation of global error between states, so that we can only say with certainty that the per-state error resulting from our method is approximately the same as with the other methods. Visual comparison of these results with the results shown in the above mentioned references provides additional evidence for this statement.

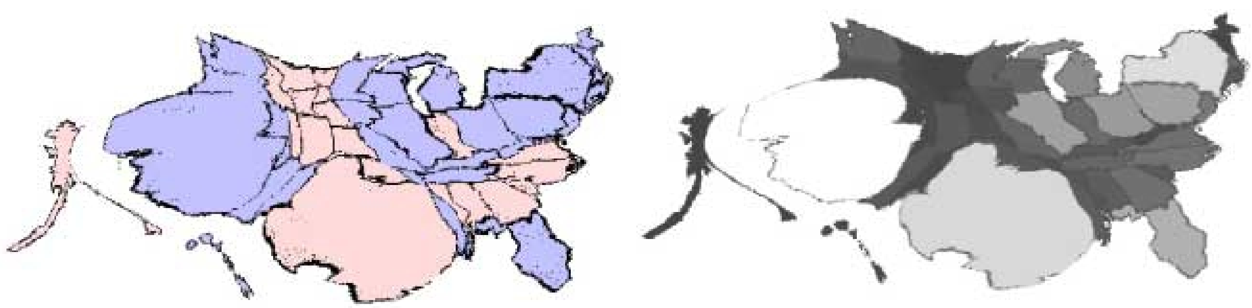

Figure 8. Normalized Views of Election Results and Electoral Votes

For performance comparisons, we ran this computation on a 200 MHz SGI workstation using a 33x33 (1069 nodes) mesh. With this setup it took 0.404 seconds of wall clock time per 100 iterations. The visual trend of the convergence is readily apparent after only 50 iterations, and by 300 iterations (1.2 seconds) the algorithm has converged to within 10% of it’s final value, as measured by pixel ratio. By 800 iterations (3.2 seconds), the algorithm has fully converged to the above result.

Example II: Interstate Speed Limits

The interstate highway system in the United States covers every state in the union, and each state is able to define the maximum speed limit on those portions of the interstates that pass through it. There is considerable variation in the speed limits chosen, from 55 miles per hour in states such as Connecticut to effectively no speed limit in Montana. All speed limits were obtained from a USENET FAQ, the numerical speed limit for Montana was arbitrarily set to a “reasonable and proper” 140 miles per hour. For a driver planning to travel across the USA, the time required for a particular route will be a function of both the geographic distance involved and the speed limits that will be enforced en-route. By encoding the speed limit information for each state as a thematic variable in a map of the USA, we can then sample that map to obtain a suitable magnification field. Here we define magnification as the inverse of the speed limit, so that states with higher speed limits will shrink to reflect the increased rate of travel. Figure 9 shows the thematic encoding of speed limits by state, along with a transformed version of the map which reflects the thematic magnification. For some of the states the transformation is quite subtle, and may not be immediately apparent when comparing static images. By animating through the changes however, the user can quite readily track these subtle changes.

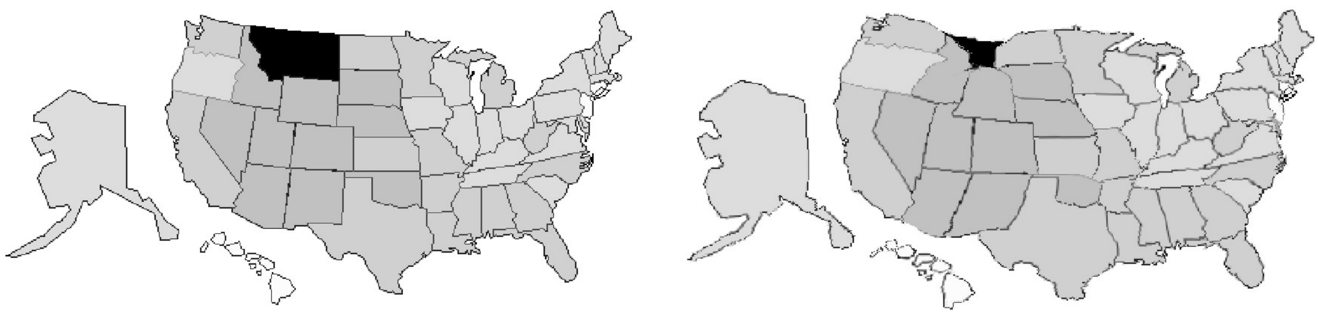

Figure 9. State Speed Limits and Normalized Driving View

Although this example is illustrative of the ability to use metrics other than quantity and density as thematic variables for this process, it also raises some questions about the ways in which the resulting views should be interpreted. When we look at the state of North Dakota (the state directly to the right of the darkest state Montana) in the normal view, we see that it is approximately rectangular in shape, so that any vertical line across the state will be approximately the same length, no matter where it is located horizontally within the state. When we look at the transformed version of the state however, we see that the left side of the state is more vertically compressed than the right side. This may give the viewer the false impression that it would be faster to drive North across the western part of the state than the eastern. The difficulty here occurs for two reasons. First is that the distribution of the thematic variable is discontinuous, while the transformation method imposes continuity on the resulting transformation. Thus the high value for Montana causes a localized influence on the surrounding regions where it interfaces with its neighbouring states. This would seem to be a problem that is inherent to all systems for producing continuous cartograms (and hence a argument in favour of discontinuous cartograms which do not enforce adjacency requirements), although [Gusein-Zade and Tikunov, 1993] claim to have had some success dealing with this type of discontinuity. The second problem is that we are overlaying a linear metric (driving distance) on a map that has been transformed using an area metric, thus the transformation method is making no effort to ensure that linear distances between regions are enforced, although on average their values will be changed suitably by the area transformation. Here again we emphasize that the intent of our method is to provide information about the *relative* impact of the thematic variable rather than the *absolute* value. Through the use of interaction and animation the user will be able to see the states grow and shrink according to their values, and thus realize the general pattern of distribution for the thematic variable, even when specific discontinuities or issues of interpretation arise.

Related Work

Comparisons between the performance of different systems is difficult, as it has not always been the common practice to report actual timing results for the convergence of the systems. Although empirical measures of the performance are not always the ultimate means of determining algorithm efficiency, time for convergence of iterative methods can be difficult to predict, and simple counts of number of iterations are often a poor indicator of overall computation time requires. We have followed the lead of [House and Kocmoud, 1998] in reporting actual time to achieve convergence, and will welcome comparison with other published figures in the future.

The method for continuous cartograms in [Tobler, 1973] is somewhat similar in nature to our iterative method, in that both systems work on regular meshes rather than on the polygonal boundaries of the region (although we note that our work was derived independently through a different set of motivating factors). Here we use texture mapping to automatically interpolate an image of the map onto the transformed space, whereas this required a distinct computational step in Tobler's system. The system described in [Dougenik et al., 1985] manages to reduce the number of iterations required for convergence while using a pre-processing step to ensure that certain constraints will not be violated. This pre-processing step shifts some computational costs away from the iterative method itself, but also introduces additional vertices to the geometry as the polygonal representation is refined. Their system was able to achieve fairly good accuracy for the US population transformation, having only 1.7% average error for each state. In [Tobler, 1986], we see a method for providing a crude but computationally inexpensive initial estimate on the transformation. The transformation can then be refined via an iterative method such as [Tobler, 1973], with the idea that the initial estimate will result in less work for the iterative method, and thus faster convergence. Similar to this, although more complex in nature is the idea of using a 2D histogram equalization technique as described in [Keahey, 1997] to provide a much more accurate initial estimate, which can then be refined using the iterative method found in that paper. The work described in [Gusein-Zade and Tikunov, 1993] is similar in motivation to [Dougenik et al., 1985] in that it makes an effort to reduce the number of iterations for convergence of their system. However, as the authors acknowledge, they are reducing the number of iterations required at the expense of greatly increasing the amount of time required for each iteration. Thus no claim has been made by them for the system running significantly faster than it's predecessors, although they do report a fairly accurate average error per-state of 1% for the US population density map. Careful attention to maintaining recognizability of regions is given in the work of [House and Kocmoud, 1998]. The authors describe an iterative method which carefully balances the constraints of area, topology and the rotation and distortion of individual regions. Their method makes use of a hierarchical polygonal representation and a simulated annealing-type process to cycle through the various constraint solvers on each iteration. They achieve fairly accurate results (1.5% average error per state) on the US population problem, for which they report a time of approximately 6 hours on a 300 MHz workstation to compute the final result.

Conclusions

Area-normalized thematic views provide a practicable method for reducing one of the most egregious «visual lies» encountered in visualization, particularly in the use of thematic maps. We have described a system which offers approximately the same accuracy as is found in other systems for continuous cartograms, while also providing convergence at near-interactive frame-rates. To the best of our knowledge, no system that converges at a rate within even a few orders of magnitude of our system has been described in the literature. By illustrating some of the shortcomings of static representations of cartograms, we emphasized the importance of interactivity in helping the user to better realize the distribution of the thematic variable.

Further Work

As mentioned in a previous section, standard multi-resolution methods would increase the value of this method greatly by reducing the number of mesh nodes required to represent large-uniform regions, while still allowing localized increases in resolution level to capture high-frequency details. Currently we provide for this with a basic multi-grid method which (based on preliminary results) has the ability to increase performance significantly, however more advanced techniques such as wavelets can take this idea much further. Such multi-resolution methods should ultimately be driven by an automated signal analysis which describes the frequency components of the map image being transformed. Very little has been done in the way of user studies on the effectiveness of continuous cartograms from the user's perspective, such a detailed study on both the static and interactive methods of presentation would likely prove beneficial to identifying shortcomings and determining whether or not further work in this area is fully warranted.

References

Blinn, J.F. and Newell, M.E. (1976). Texture and reflection in computer-generated images. Communications of the ACM, 19(10).

Churcher, N., Prachuabmoh, P. and Churcher, C. (1997). Visualization techniques for collaborative GIS browsers. In International Conference on GeoComputation.

Dougenik, J.A., Chrisman, N.R. and Niemeyer, D.R. (1985). An algorithm to construct continuous area cartograms. Professional Geographer, 37(1).

Gusein-Zade, S.M. and Tikunov, V.S. (1993). A new technique for constructing continuous cartograms. Cartography and Geographic Information Systems, 20(3).

House, D.H. and Kocmoud, C.J. (1998). Continuous cartogram construction. IEEE Visualization.

Kadmon, N. and Shlomi, E. (1978). A polyfocal projection for statistical surfaces. Cartographic Journal, 15(1):36-41.

Keahey, T.A. and Robertson, E.L. (1996). Techniques for non-linear magnification transformations. IEEE Symposium on Information Visualization.

Keahey, T.A. and Robertson, E.L. (1997). Nonlinear magnification fields. IEEE Symposium on Information Visualization.

Keahey, T.A. (1997). Nonlinear Magnification. PhD thesis, Department of Computer Science, Indiana University.

Leung, Y.K. and Apperley, M.D. (1994). A review and taxonomy of distortion-oriented presentation techniques. ACM Transactions on Computer-Human Interaction, 1(2):126-160.

Munzner, T. (1997). H3: laying out large directed graphs in 3D hyperbolic space. IEEE Symposium on Information Visualization.

Schaffer, D., Zuo, Z., Bartram, L., Dill, J., Dubs, S., Greenberg, S., Roseman, M. (1993). Comparing fisheye and full-zoom techniques for navigation of hierarchically clustered networks. Graphics Interface.

Tobler, W.R. (1973). A continuous transformation useful for districting. Annals, New York Academy of Science, 219.

Tobler, W.R. (1986). Pseudo-cartograms. The American Cartographer, 13(1).

Tufte, E.R. (1983). The visual display of quantitative info